

A correspondence-less approach to matching of deformable shapes

Jonathan Pokrass¹, Alexander M. Bronstein¹, and Michael M. Bronstein²

¹Dept. of Electrical Engineering, Tel Aviv University, Israel
 evgenyfo@post.tau.ac.il, bron@eng.tau.ac.il

²Inst. of Computational Science, Faculty of Informatics,
 Università della Svizzera Italiana, Lugano, Switzerland
 michael.bronstein@usi.ch

Abstract. Finding a match between partially available deformable shapes is a challenging problem with numerous applications. The problem is usually approached by computing local descriptors on a pair of shapes and then establishing a point-wise correspondence between the two. In this paper, we introduce an alternative correspondence-less approach to matching fragments to an entire shape undergoing a non-rigid deformation. We use diffusion geometric descriptors and optimize over the integration domains on which the integral descriptors of the two parts match. The problem is regularized using the Mumford-Shah functional. We show an efficient discretization based on the Ambrosio-Tortorelli approximation generalized to triangular meshes. Experiments demonstrating the success of the proposed method are presented.

Keywords: deformable shapes, partial matching, partial correspondence, partial similarity, diffusion geometry, Laplace-Beltrami operator, shape descriptors, heat kernel signature, Mumford-Shah regularization

1 Introduction

In many real-world settings of the shape recognition problem, the data are degraded by acquisition imperfections and noise, resulting in the need to find *partial* similarity of objects. Such cases are common, for example, in face recognition, where the facial surface may be partially occluded by hair. In other applications, such as shape retrieval, correct semantic similarity of two objects is based on partial similarity – for example, a centaur is partially similar to a human because they share the human-like upper body [15].

In rigid shape analysis, modifications of the popular iterative closest point (ICP) algorithm are able to deal with partial shape alignment by rejecting points with bad correspondences (e.g., by thresholding the product of local normal vectors). However, it is impossible to guarantee how large and regular the resulting corresponding parts will be.

Bronstein *et al.* [4] formulated non-rigid partial similarity as a multi-criterion optimization problem, in which one tries to find the corresponding parts in two

shapes by simultaneously maximizing significance and similarity criteria (in [4], metric distortion [14,19,5] was used as a criterion of similarity, and part area as significance). The problem requires the knowledge of correspondence between the shapes, and in the absence of a given correspondence, can be solved by alternating between weighted correspondence finding and maximization of part area. In [4], a different significance criterion based on statistical occurrence of local shape descriptors was used.

One of the drawbacks of the above method is its tendency in some cases to find a large number of disconnected components, which have the same area as a larger single component. The same authors addressed this problem using a Mumford-Shah [22,10]-like regularization for rigid [3] and non-rigid [2] shapes.

Recent works on local shape descriptors (see, e.g., descriptors [16,33,24,11,18,21,32,29,27,9]) have led to the adoption of bags of features [28] approach popular in image analysis for the description of 3D shapes [21,23,30]. Bags of features allows to some extent finding partial similarity, if the overlap between the parts is sufficiently large.

In this paper, we present an approach for correspondence-less partial matching of non-rigid 3D shapes. Our work is inspired by the recent work on partial matching of images [12]. The main idea of this approach, adopted here, is to find similar parts by comparing part-wise distributions of local descriptors. This removes the need of correspondence knowledge and greatly simplifies the problem.

The rest of the paper is organized as follows. In Section 2, we review the mathematical background of diffusion geometry, which is used for the construction of local descriptors. Section 3 deals with the partial matching problem and Section 4 addresses its discretization. Section 5 presents experimental results. Finally, Section 6 concludes the paper.

2 Background

Diffusion geometry. Diffusion geometry is an umbrella term referring to geometric analysis of diffusion or random walk processes. We model a shape as a compact two-dimensional Riemannian manifold X . In its simplest setting, a diffusion process on X is described by the partial differential equation

$$\left(\frac{\partial}{\partial t} + \Delta\right) f(t, x) = 0, \quad (1)$$

called the *heat equation*, where Δ denotes the positive-semidefinite Laplace-Beltrami operator associated with the Riemannian metric of X . The heat equation describes the propagation of heat on the surface and its solution $f(t, x)$ is the heat distribution at a point x in time t . The initial condition of the equation is some initial heat distribution $f(0, x)$; if X has a boundary, appropriate boundary conditions must be added.

The solution of (1) corresponding to a point initial condition $f(0, x) = \delta(x, y)$, is called the *heat kernel* and represents the amount of heat transferred from x to y in time t due to the diffusion process. The value of the heat kernel $h_t(x, y)$

can also be interpreted as the transition probability density of a random walk of length t from the point x to the point y .

Using spectral decomposition, the heat kernel can be represented as

$$h_t(x, y) = \sum_{i \geq 0} e^{-\lambda_i t} \phi_i(x) \phi_i(y). \quad (2)$$

Here, ϕ_i and λ_i denote, respectively, the eigenfunctions and eigenvalues of the Laplace-Beltrami operator satisfying $\Delta \phi_i = \lambda_i \phi_i$ (without loss of generality, we assume λ_i to be sorted in increasing order starting with $\lambda_0 = 0$). Since the Laplace-Beltrami operator is an *intrinsic* geometric quantity, i.e., it can be expressed solely in terms of the metric of X , its eigenfunctions and eigenvalues as well as the heat kernel are invariant under isometric transformations (bending) of the shape.

Heat kernel signatures. By setting $y = x$, the heat kernel $h_t(x, x)$ expresses the probability density of remaining at a point x after time t . The value $h_t(x, x)$, sometimes referred to as the *auto-diffusivity function*, is related to the Gaussian curvature $K(x)$ through

$$h_t(x, x) \approx \frac{1}{4\pi t} \left(1 + \frac{1}{6} K(x) t + \mathcal{O}(t^2) \right). \quad (3)$$

This relation coincides with the well-known fact that heat tends to diffuse slower at points with positive curvature, and faster at points with negative curvature. Under mild technical conditions, the set $\{h_t(x, x)\}_{t>0}$ is fully informative in the sense that it allows to reconstruct the Riemannian metric of the manifold [29].

Sun *et al.* [29] proposed constructing point-wise descriptors referred to as *heat kernel signatures* (HKS) by taking the values of the discrete auto-diffusivity function at point x at multiple times, $\mathbf{p}(x) = c(x)(h_{t_1}(x, x), \dots, h_{t_d}(x, x))$, where t_1, \dots, t_d are some fixed time values and $c(x)$ is chosen so that $\|p(x)\|_2 = 1$. Such a descriptor is a vector of dimensionality d at each point. Since the heat kernel is an intrinsic quantity, the HKS is invariant to isometric transformations of the shape.

A scale-invariant version of the HKS descriptor (SI-HKS) was proposed in [9]. First, the heat kernel is sampled logarithmically in time. Next, the logarithm and a derivative with respect to time of the heat kernel values are taken to undo the multiplicative constant. Finally, taking the magnitude of the Fourier transform allows to undo the scaling of the time variable.

Bags of features. Ovsjanikov *et al.* [23] and Toldo *et al.* [30] proposed constructing global shape descriptors from local descriptors using the *bag of features* paradigm [28]. In this approach, a fixed “geometric vocabulary” is computed by means of an off-line clustering of the descriptor space. Next, each point descriptor is represented in the vocabulary using vector quantization. The bag of features global shape descriptor is then computed as the histogram of quantized descriptors over the entire shape.

3 Partial matching

In what follows, we assume to be given two shapes X and Y with corresponding point-wise descriptor fields \mathbf{p} and \mathbf{q} defined on them (here we adopt HKS descriptors, though their quantized variants or any other intrinsic point-wise descriptors can be used as well). Assuming that Y is a part of an unknown shape that is intrinsically similar to X , we aim at finding a part $X' \subseteq X$ having the same area A of Y such that the integral shape descriptors computed on X' and Y coincide as closely as possible. In order to prevent the parts from being fragmented and irregular, we penalize for their boundary length. The entire problem can be expressed as minimization of the following energy functional

$$E(X') = \left\| \int_{X'} \mathbf{p} da - \bar{\mathbf{q}} \right\|^2 + \lambda_r L(\partial X') \quad (4)$$

under the constraint $A(X') = A$, where A denotes area and $\bar{\mathbf{q}} = \int_Y \mathbf{q} da$. The first term of the functional constitutes the data term while the second one is the regularity term whose influence is controlled by the parameter λ_r .

Discretization of the above minimization problem with a crisp set X' results in combinatorial complexity. To circumvent this difficulty, in [2,3] it was proposed to relax the problem by replacing the crisp part X' by a fuzzy membership function u on X , replacing the functional E by a generalization of the Mumford-Shah functional [22] to surfaces. Here, we adopt this relaxation as well as the approximation of the Mumford-Shah functional proposed by Ambrosio and Tortorelli [1]. This yields the problem of the form

$$\begin{aligned} \min_{u, \rho, \sigma} D(u) + \lambda_r R(u; \rho) \\ \text{s.t. } \int_X u da = A, \end{aligned} \quad (5)$$

with the data term

$$D(u) = \left\| \int_X \mathbf{p} u da - \bar{\mathbf{q}} \right\|^2 \quad (6)$$

and the Ambrosio-Tortorelli regularity term

$$R(u; \rho) = \frac{\lambda_s}{2} \int_X \rho^2 \|\nabla u\|^2 da + \lambda_b \epsilon \int_X \|\nabla \rho\|^2 da + \frac{\lambda_b}{4\epsilon} \int_X (1 - \rho)^2 da, \quad (7)$$

where ρ is the so-called phase field indicating the discontinuities of u , and $\epsilon > 0$ is a parameter.

The first term of R above imposes piece-wise smoothness of the fuzzy part u governed by the parameter λ_s . By setting a sufficiently large λ_s , the parts become approximately piece-wise constant as desired in the original crisp formulation (4). The second term of R is analogous to the boundary length term in (4) and converges to the latter as $\epsilon \rightarrow 0$.

We minimize (5) using alternating minimization comprising the following two iteratively repeated steps:

Step 1: fix ρ and solve for u

$$\min_u \left\| \int_X \mathbf{p} u da - \bar{\mathbf{q}} \right\|^2 + \lambda_r \frac{\lambda_s}{2} \int_X \rho^2 \|\nabla u\|^2 da \quad \text{s.t.} \quad \int_X u da = A. \quad (8)$$

Step 2: fix the part u and solve for ρ

$$\min_{\rho} \frac{\lambda_s}{2} \int_X \rho^2 \|\nabla u\|^2 da + \lambda_b \epsilon \int_X \|\nabla \rho\|^2 da + \frac{\lambda_b}{4\epsilon} \int_X (1 - \rho)^2 da. \quad (9)$$

4 Discretization and numerical aspects

We represent the surface X as triangular mesh with n faces constructed upon the samples $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ and denote by $\mathbf{a} = (a_1, \dots, a_m)^T$ the corresponding area elements at each vertex. $\mathbf{A} = \text{diag}\{\mathbf{a}\}$ denote the diagonal $m \times m$ matrix created out of \mathbf{a} . The membership function u is sampled at each vertex and represented as the vector $\mathbf{u} = (u_1, \dots, u_m)^T$. Similarly, the phase field is represented as the vector $\boldsymbol{\rho} = (\rho_1, \dots, \rho_m)^T$.

Descriptors. The computation of the discrete heat kernel $h_t(\mathbf{x}_1, \mathbf{x}_2)$ requires computing discrete eigenvalues and eigenfunctions of the discrete Laplace-Beltrami operator. The latter can be computed directly using the finite elements method (FEM) [26], or by discretization of the Laplace operator on the mesh followed by its eigendecomposition. Here, we adopt the second approach according to which the discrete Laplace-Beltrami operator is expressed in the following generic form,

$$(\Delta f)_i = \frac{1}{a_i} \sum_j w_{ij} (f_i - f_j), \quad (10)$$

where $f_i = f(\mathbf{x}_i)$ is a scalar function defined on the mesh, w_{ij} are weights, and a_i are normalization coefficients. In matrix notation, (10) can be written as $\Delta \mathbf{f} = \mathbf{A}^{-1} \mathbf{W} \mathbf{f}$, where \mathbf{f} is an $m \times 1$ vector and $\mathbf{W} = \text{diag} \left\{ \sum_{l \neq i} w_{il} \right\} - w_{ij}$.

The discrete eigenfunctions and eigenvalues are found by solving the *generalized eigendecomposition* [17] $\mathbf{W} \boldsymbol{\Phi} = \mathbf{A} \boldsymbol{\Phi} \boldsymbol{\Lambda}$, where $\boldsymbol{\Lambda} = \text{diag}\{\lambda_l\}$ is a diagonal matrix of eigenvalues and $\boldsymbol{\Phi} = (\phi_l(x_i))$ is the matrix of the corresponding eigenvectors.

Different choices of \mathbf{W} have been studied, depending on which continuous properties of the Laplace-Beltrami operator one wishes to preserve [13,31]. For triangular meshes, a popular choice adopted in this paper is the *cotangent weight* scheme [25,20], in which

$$w_{ij} = \begin{cases} (\cot \beta_{ij} + \cot \gamma_{ij})/2 : \mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i); \\ 0 : \text{else}, \end{cases} \quad (11)$$

where β_{ij} and γ_{ij} are the two angles opposite to the edge between vertices \mathbf{x}_i and \mathbf{x}_j in the two triangles sharing the edge.

Data term. Denoting by \mathbf{P} $d \times m$ the matrix of point-wise descriptors on X (stored in columns), we have

$$\int \mathbf{p} u da \approx \mathbf{P} \text{diag}\{\mathbf{A}\} \mathbf{u} \quad (12)$$

This yields the following discretization of the data term (6)

$$D(\mathbf{u}) = \|\mathbf{P}\mathbf{A}\mathbf{u} - \bar{\mathbf{q}}\|^2 = \mathbf{u}^T \mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{A} \mathbf{u} - 2\bar{\mathbf{q}}^T \mathbf{P} \mathbf{A} \mathbf{u} + \bar{\mathbf{q}}^T \bar{\mathbf{q}}. \quad (13)$$

Gradient norm. We start by deriving the discretization of a single term $\rho^2 \|\nabla u\|^2 da$ in some triangle t of the mesh. Let us denote by $\mathbf{x}_i, \mathbf{x}_j$ and \mathbf{x}_k the vertices of the triangle and let $\mathbf{X}_t = (\mathbf{x}_j - \mathbf{x}_i, \mathbf{x}_k - \mathbf{x}_i)$ be the 3×2 matrix whose columns are the vectors forming the triangle, and by $\alpha_t = \frac{1}{2} \sqrt{\det(\mathbf{X}_t^T \mathbf{X}_t)}$ its area. Let also \mathbf{D}_t be the sparse $2 \times m$ matrix with $+1$ at indices $(1, j)$ and $(2, k)$, and -1 at $(1, i)$ and $(2, i)$. \mathbf{D}_t is constructed in such a way to give the differences of values of u on the vertices of the triangle with respect to the values at the central vertex, $\mathbf{D}_t \mathbf{u} = (u_j - u_i, u_k - u_i)^T$. The gradient of the function u is constant on the triangle and can be expressed in these terms by $\mathbf{g}_t = (\mathbf{X}_t^T \mathbf{X}_t)^{-1/2} \mathbf{D}_t \mathbf{u} = \mathbf{G}_t \mathbf{u}$.

In order to introduce the weighting by ρ^2 , let \mathbf{S} be an $n \times m$ sparse matrix with the elements $s_{ti} = \frac{1}{3}$ for every vertex i belonging to the triangle t and zero otherwise. In this notation, $\mathbf{S}\boldsymbol{\rho}$ is a per-triangle field whose elements are the average values of $\boldsymbol{\rho}$ on each of the mesh triangles. We use the Kronecker product of \mathbf{S} with $\mathbf{1} = (1, 1)^T$ to define the $2n \times n$ matrix $\mathbf{S} \otimes \mathbf{1}$ formed by replicating twice each of the rows of \mathbf{S} . This yields

$$\begin{aligned} \int \rho^2 \|\nabla u\|^2 da &\approx \sum_t \|(\mathbf{S}\boldsymbol{\rho})_t \mathbf{G}_t \mathbf{u}\|^2 \alpha_t = \\ &= \left\| \text{diag}\{(\mathbf{S} \otimes \mathbf{1})\boldsymbol{\rho}\} \begin{pmatrix} \sqrt{\alpha_1} \mathbf{G}_1 \\ \vdots \\ \sqrt{\alpha_n} \mathbf{G}_n \end{pmatrix} \right\|^2 = \mathbf{u}^T \mathbf{G}^T \mathbf{S}(\boldsymbol{\rho}) \mathbf{G} \mathbf{u}, \end{aligned} \quad (14)$$

where \mathbf{G} is the matrix containing $\sqrt{\alpha_t} \mathbf{G}_t$ stacked as rows, and $\mathbf{S}(\boldsymbol{\rho}) = \text{diag}\{(\mathbf{S} \otimes \mathbf{1})\boldsymbol{\rho}\}^2$.

Discretized alternating minimization. We plug in the results obtained so far into the two steps of the alternating minimization problem (8)–(9). For fixed $\boldsymbol{\rho}$, the discretized minimization problem (8) w.r.t. \mathbf{u} can be written as

$$\min_{\mathbf{u}} \mathbf{u}^T \left(\mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{A} + \lambda_r \frac{\lambda_s}{2} \mathbf{G}^T \mathbf{S}(\boldsymbol{\rho}) \mathbf{G} \right) \mathbf{u} - 2\bar{\mathbf{q}}^T \mathbf{P} \mathbf{A} \mathbf{u} \quad \text{s.t.} \quad \mathbf{a}^T \mathbf{u} = A \quad (15)$$

Let us now fix \mathbf{u} . In a triangle t , we denote $g_t^2 = \|\mathbf{G}_t \mathbf{u}\|^2$ and let $\mathbf{R}(\mathbf{u}) = \text{diag}\{\alpha_1 g_1^2, \dots, \alpha_n g_n^2\}$. Using this notation, we obtain the following discretization of the integrals in the regularization term (7)

$$\int_X \rho^2 \|\nabla u\|^2 da \approx \sum_t (\mathbf{S}\boldsymbol{\rho})_t^2 g_t^2 \alpha_t = \boldsymbol{\rho}^T \mathbf{S}^T \mathbf{R}(\mathbf{u}) \mathbf{S} \boldsymbol{\rho}. \quad (16)$$

Similar to the derivation of (14),

$$\int \|\nabla \rho\|^2 da \approx \boldsymbol{\rho}^T \mathbf{S} \mathbf{G}^T \mathbf{G} \mathbf{S} \boldsymbol{\rho}, \quad (17)$$

and

$$\int_X (1 - \rho)^2 da \approx \boldsymbol{\rho}^T \mathbf{A} \boldsymbol{\rho} - 2 \mathbf{a}^T \boldsymbol{\rho} + \mathbf{1}^T \mathbf{a}. \quad (18)$$

The discretized minimization problem (9) w.r.t. $\boldsymbol{\rho}$ becomes

$$\min_{\boldsymbol{\rho}} \boldsymbol{\rho}^T \left(2\epsilon \lambda_s \mathbf{S}^T \mathbf{R}(\mathbf{u}) \mathbf{S} + 4\epsilon^2 \lambda_b \mathbf{S} \mathbf{G}^T \mathbf{G} \mathbf{S} + \lambda_b \mathbf{A} \right) \boldsymbol{\rho} - 2\lambda_b \mathbf{a}^T \boldsymbol{\rho}. \quad (19)$$

Since the above is an unconstrained quadratic problem, it has the following closed-form solution

$$\boldsymbol{\rho} = \left(\frac{2\epsilon \lambda_s}{\lambda_b} \mathbf{S}^T \mathbf{R}(\mathbf{u}) \mathbf{S} + 4\epsilon^2 \mathbf{S} \mathbf{G}^T \mathbf{G} \mathbf{S} + \mathbf{A} \right)^{-1} \mathbf{a}. \quad (20)$$

5 Results

In order to test our approach, we performed several partial matching experiments on data from the SHREC 2010 benchmark [8,7] and the TOSCA dataset [6].¹ The datasets contained high-resolution (10K-50K vertices) triangular meshes of human and animal shapes with simulated transformations, with known groundtruth correspondence between the transformed shapes. In our experiments, all the shapes were downsampled to approximately 2500 vertices. Parts were cut by taking a geodesic circle of random radius around a random center point.

For each part, the normalized HKS descriptor was calculated at each vertex belonging to the part. To avoid boundary effects (see Figure 1), descriptors close to the boundary were ignored when calculating $\bar{\mathbf{q}}$ in (4). The distance from the boundary was selected in accordance to the time scales of the HKS. We used ten linearly spread samples in range [65, 90] for the descriptors and the according distance taken from edge was set to 15. Two to three iterations of the alternating minimization procedure were used, exhibiting fast convergence (Figure 2). After three iterations the member function \mathbf{u} typically ceased changing significantly. The phase map ρ assumed the values close to 1 in places of low gradient of the membership function \mathbf{u} , and less than 1 in high gradient areas (Figure 2). The importance of the regularization step is evident observing the change in \mathbf{u} in Figure 2. Figure 3 shows the influence of the parameter λ_r , controlling the impact of the regularization. For too small values of λ_r , two equally weighted matches are obtained due to symmetry (left). The phenomenon decreases with the increase of the influence of the regularization penalty. However, increasing λ_r more causes incorrect matching (second and third columns from the right) due

¹ Both datasets are available online at <http://tosca.cs.technion.ac.il>

to low data term influence. Increasing it even more starts smoothing the result (rightmost column) until eventually making the membership function uniform over the entire shape. The resulting membership function \mathbf{u} was thresholded in such a way that the outcome area will be as close as possible to the query area. Figures 4–6 show examples of matching results after thresholding. Notice that in Figure 4 the matching result sometimes contain the symmetric counter part of the result due to invariance of the HKS descriptor to intrinsic symmetry. (in this figure, the threshold was adjusted to the value of 0.35 when the membership functions weights are split between two symmetric parts as in Figure 3). The method is robust to shape deformations and geometric and topological noise as depicted in Figure 6. Note that the figures show part-to-whole shape matching, but because of the low scale HKS descriptors the same procedure works for matching to other parts as well. Table 1 summarizes quantitative evaluation




Fig. 1. An RGB visualization of the first three component of HKS descriptors computed on the full shape (left) and on a part of the shape (center). The L_2 difference between the two fields is depicted in the rightmost figure. Note that the difference is maximal on the boundary decaying fast away from it; the error decay speed depends on the scale choice in the HKS descriptor.

that was performed on a subset of the SHREC database, for which groundtruth correspondence and its bilaterally symmetric counterpart were available. This subset included a male, a dog and a horse shape classes with different geometric, topological and noise deformations (98 shapes in total). The query set was generated by selecting a part from a deformed shape (1000 queries in each deformation category) and matched to the null shape with parameters and thresholds as described above.

Complexity. The code was implemented in Matlab with some parts written in C with MEX interface. The quadratic programming problem (8) in Step 1 was solved using QPC² implementation of a dual active set method. The experiments were run on 2.3GHz Intel Core2 Quad CPU, 2GB RAM in Win7 32bit environ-

² available online at <http://sigpromu.org/quadprog>




Fig. 2. Convergence of the alternating minimization procedure. Depicted are the membership function \mathbf{u} (top row) and the phase field ρ (bottom row) at the first three iterations.




Fig. 3. The influence of the parameter λ_r , controlling the impact of regularization. The leftmost figure depicts a query part; figures on the right are the membership function u for different values of λ_r .

ment. The running time (including re-meshing and descriptor calculation) per part was 40 – 50 sec.

6 Conclusions

We presented a framework for finding partial similarity between shapes which does not rely on explicit correspondence. The method is based on regularized matching of region-wise local descriptors, and can be efficiently implemented. Experimental results show that our approach performs well in challenging matching scenarios, such as the presence of geometric and topological noise. In the future work, we will extend the method to the setting of two partially-similar full shapes, in which two similar parts have to be found in each shape, and then consider a multi-part matching (puzzle) scenario.

Table 1. Part matching performance on transformed shapes from the SHREC benchmark. At each query a random part (location and size) was selected from a deformed shape and matched to the null shape. Overlap is reported compared to the groundtruth correspondence between the shapes (in parentheses taking into account the intrinsic bilateral symmetry).

Transformation	Queries	Avg. overlap
Isometry	1000	75% (85%)
Isometry + Shotnoise	1000	75% (85%)
Isometry + Noise	1000	71% (82%)
Isometry + Microholes	1000	68% (82%)
Isometry + Holes	1000	66% (76%)
All	5000	71% (82%)

References

1. L. Ambrosio and V.M. Tortorelli. Approximation of functionals depending on jumps by elliptic functionals via-convergence. *Comm. Pure Appl. Math*, 43(8):999–1036, 1990.
2. A. M. Bronstein and M. M. Bronstein. Not only size matters: regularized partial matching of nonrigid shapes. In *Prof. NORDIA*, 2008.
3. A. M. Bronstein and M. M. Bronstein. Regularized partial matching of rigid shapes. In *Proc. ECCV*, pages 143–154, 2008.
4. A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Partial similarity of objects, or how to compare a centaur to a horse. *IJCV*, 84(2):163–183, 2009.
5. A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Science (PNAS)*, 103(5):1168–1172, 2006.
6. A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer-Verlag New York Inc, 2008.
7. AM Bronstein, MM Bronstein, B. Bustos, U. Castellani, M. Crisani, B. Falcidieno, LJ Guibas, I. Kokkinos, V. Murino, M. Ovsjanikov, et al. SHREC 2010: robust feature detection and description benchmark. *Proc. 3DOR*, 2010.
8. AM Bronstein, MM Bronstein, U. Castellani, A. Dubrovina, LJ Guibas, RP Horeaud, R. Kimmel, D. Knossow, E. von Lavante, D. Mateus, et al. SHREC 2010: robust correspondence benchmark. In *Eurographics Workshop on 3D Object Retrieval*, 2010.
9. M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*, 2010.
10. T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, 2001.
11. U. Clarenz, M. Rumpf, and A. Telea. Robust feature detection and local classification for surfaces based on moment analysis. *Trans. Visualization and Computer Graphics*, 10(5):516–524, 2004.
12. C. Domokos and Z. Kato. Affine Puzzle: Realigning Deformed Object Fragments without Correspondences. In *Proc. ECCV*, pages 777–790, 2010.

13. M.S. Floater and K. Hormann. Surface parameterization: a tutorial and survey. *Advances in Multiresolution for Geometric Modelling*, 1, 2005.
14. M. Gromov. *Structures Métriques Pour les Variétés Riemanniennes*. Number 1 in Textes Mathématiques. 1981.
15. D. Jacobs, D. Weinshall, and Y. Gdalyahu. Class representation and image retrieval with non-metric distances. *Trans. PAMI*, 22(6):583–600, 2000.
16. A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3D scenes. *Trans. PAMI*, 21(5):433–449, 1999.
17. B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Proc. Shape Modeling and Applications*, 2006.
18. S. Manay, B.W. Hong, A.J. Yezzi, and S. Soatto. Integral invariant signatures. *Lecture Notes in Computer Science*, pages 87–99, 2004.
19. F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5:313–346, 2005.
20. M. Meyer, M. Desbrun, P. Schroder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics III*, pages 35–57, 2003.
21. N. J. Mitra, L. J. Guibas, J. Giesen, and M. Pauly. Probabilistic fingerprints for shapes. In *Proc. SGP*, 2006.
22. D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on pure and applied mathematics*, 42(5):577–685, 1989.
23. M. Ovsjanikov, A.M. Bronstein, L.J. Guibas, and M.M. Bronstein. Shape Google: a computer vision approach to invariant shape retrieval. In *Proc. NORDIA*. Citeseer, 2009.
24. M. Pauly, R. Keiser, and M. Gross. Multi-scale feature extraction on point-sampled surfaces. In *Computer Graphics Forum*, volume 22, pages 281–289, 2003.
25. U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993.
26. Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. Laplace-spectra as fingerprints for shape matching. In *Proc. ACM Symp. Solid and Physical Modeling*, pages 101–106, 2005.
27. I. Sipiran and B. Bustos. A robust 3D interest points detector based on Harris operator. In *Proc. 3DOR*, pages 7–14. Eurographics, 2010.
28. J. Sivic and A. Zisserman. Video Google: a text retrieval approach to object matching in videos. In *Proc. CVPR*, 2003.
29. J. Sun, M. Ovsjanikov, and L. Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer Graphics Forum*, volume 28, pages 1383–1392, 2009.
30. R. Toldo, U. Castellani, and A. Fusiello. Visual vocabulary signature for 3D object retrieval and partial matching. In *Proc. 3DOR*, 2009.
31. M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: no free lunch. In *Conf. Computer Graphics and Interactive Techniques*, 2008.
32. A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009.
33. C. Zhang and T. Chen. Efficient feature extraction for 2D/3D objects in mesh representation. In *Proc. ICIP*, volume 3, 2001.




Fig. 4. Examples of matching of random parts of shapes (first row) to approximately isometric deformations of the shapes (second row). Color code indicates different parts.




Fig. 5. Examples of matching of random parts of shapes (first row) to to approximately isometric deformations of the shapes (second row).



Fig. 6. Results of matching a part of shape (first row) to shapes distorted by different transformations (second row). Shown left-to-right are: shot noise, noise, micro holes and holes.